OCTOBER 1999

SMPTE Journal

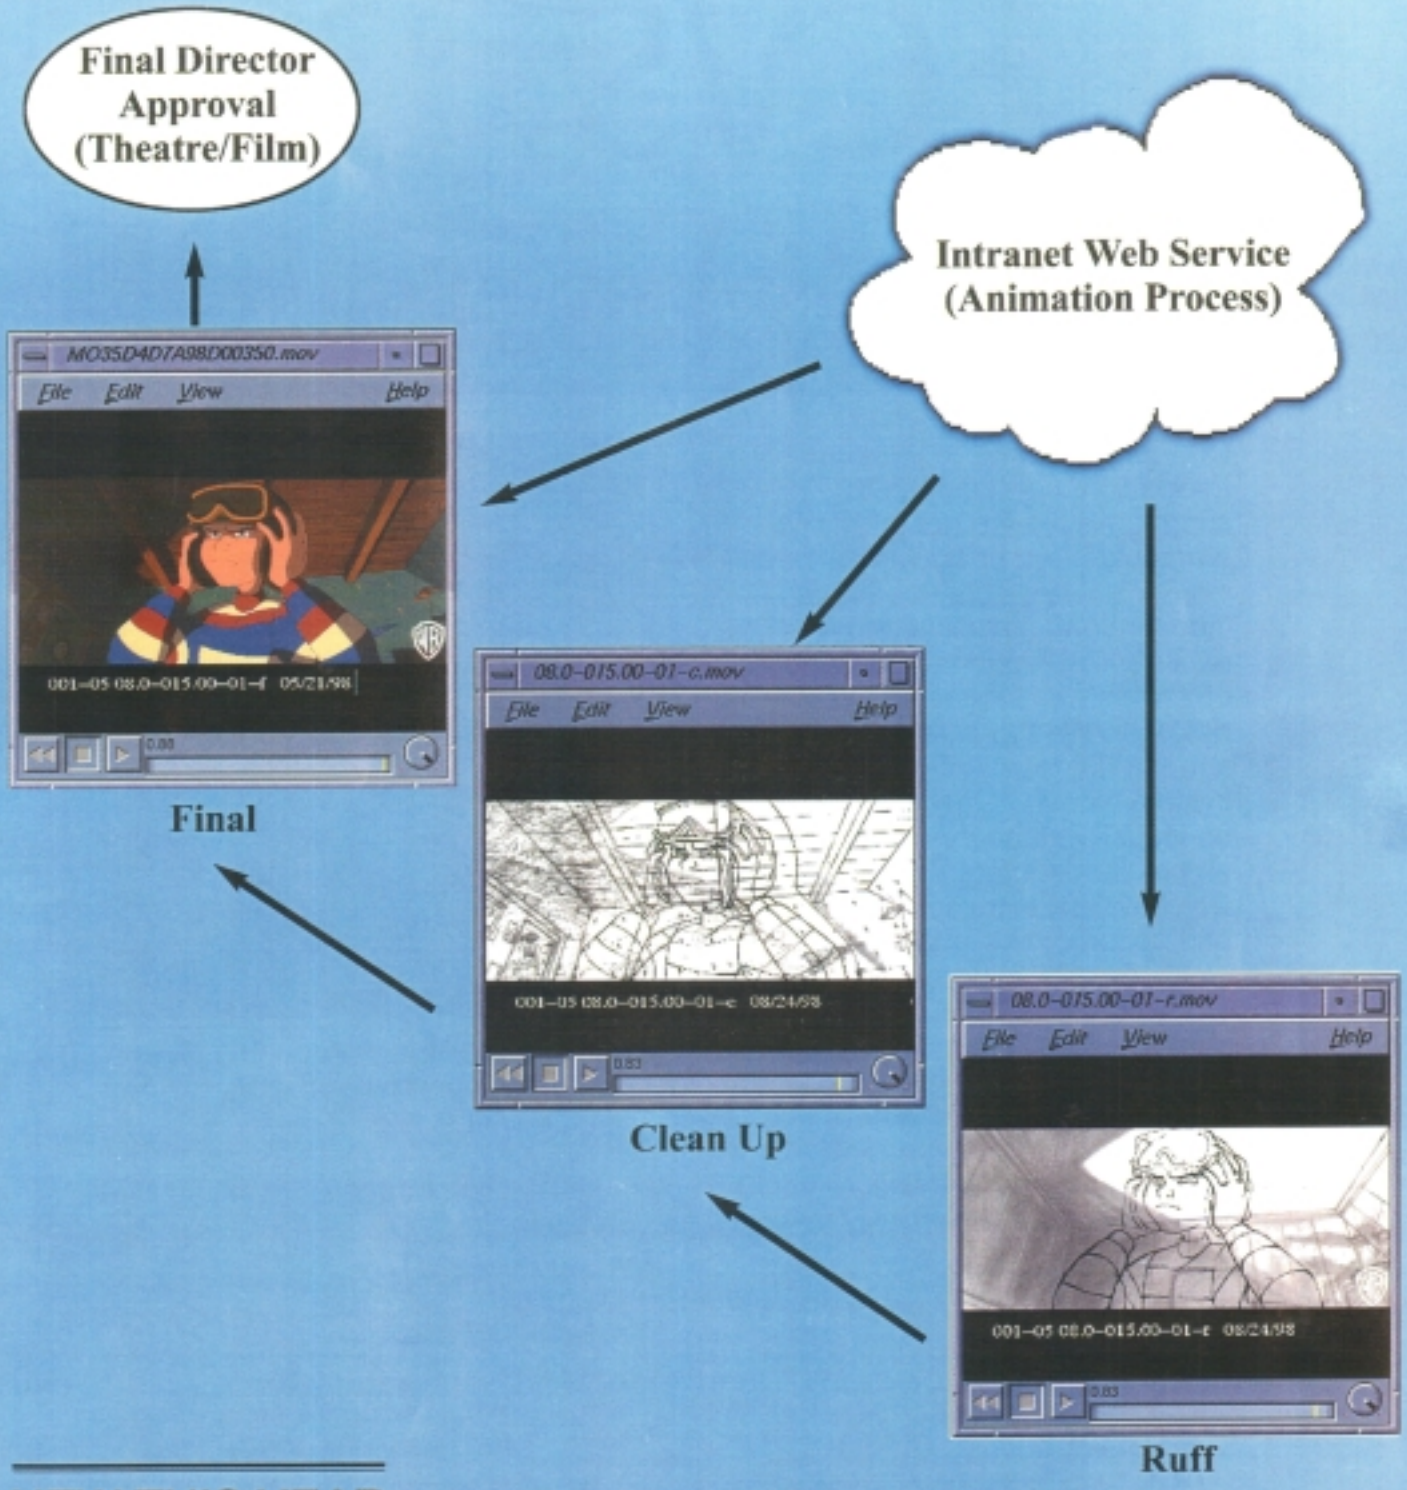

Implementation of Intranet Scene Preview for Feature Animation

By Leonard J. Reder and Gene Takahashi

This paper describes the development of an Intranet scene preview capability within Warner Bros. Feature Animation. This capability was developed so all members of the production staff could easily view production material as it was being animated. The animation production process as it pertains to the preview capabilities is presented and the organization of the server is reviewed. A discussion of HyperText MarkUp Language (HTML) along with the structure and function of the custom software developed is given. Peerformance and enhancements of QuickTime downloading and playback is also discussed.

Feature animated motion pictures typically require the work of hundreds of people to create the final product. The process, involving many specialized departments, consists of creating a story reel from sketches representing the film and then generating mechanical tests (for example, camera moves), ruff, clean up, and final (color) animations. The film to be animated is divided into sequences and then into scenes. Each scene moves from department to department at various rates depending on complexity. Rendered animation for each scene is stored on central file servers in the form of high-resolution sequences of TIFF image files and other associated data elements.

During the course of a production, scenes must be archived to tape due to disk space limitations. Since the animation is stored as high-resolution digital material, it cannot be viewed by members of the crew who only have access to desktop computers. Many times material is not online when production personnel need to view "hookups" of scenes (that is, the scene they are working on in conjunction with the previous and following scenes).

A mechanism for maintaining scenes from each production online in a form accessible to any type of computer was needed. The solution was to utilize the company's Intranet allowing animated material to be viewed as it exists at any time during the production of an animated feature.

Animated scene content was served for preview capability via an Intranet web server. Software was developed to automatically generate QuickTime movies and to create HTML pages for accessing them. Annotation was placed on every frame of each movie, and then each movie clip was striped with production audio. The software ran automatically every night and made updates to the web pages while generating new QuickTime files for scenes that had changed. For example, a scene that consisted of ruff animation and then changed to clean up would be updated. Also, if a scene was detected to have had an increase in take number because of changes requested by the director, it would also be updated. An e-mail message notifying users of exactly which scenes had been updated was also generated.

Two types of QuickTime movies were served for preview: low-resolution (320 x 240 pixel frames) compressed with the Apple video codec and 720 x 540 pixel JPEG-compressed. The 320 x 240 frame size made it possible for any modem computer system to load movies in a reasonable amount of time and to play them at, or near, 24 frame/sec film playback rate. The 720 x 540 format is D-1 video compatible resolution used with desktop nonlinear editing systems.

An editorial Intranet website has been configured to serve the scene preview QuickTime movies. The server also contains low-resolution QuickTime story reel movies uploaded from the editorial department. The movies are exported from the Avid Film Composer nonlinear edit system. These QuickTimes allow anyone to refer to the most current version of the story reel from any computer connected to the company's Intranet. Documents that describe procedures and naming convention are also available via the web, and most recently, sound effects retrieval has been added.

Animation Production Process1

After a story reel of sketches has been locked and approved by the director, production dialog is recorded and added to it in place of temporary dialog; once this is completed the animation process starts. First, the story reel is broken down into 2 to 5-min sequences and further into individual scenes (for example, shots) that can be animated. These scenes go through workbook and layout where each element is further defined and the action is planned in detail (Fig. 1). Production dialog track is "read" and sent to a central file server for use by animators creating drawings. The track reading process defines the phoneme associated with each frame of animation.

Ruff animation is a sequence of drawings created (based on the layout, workbook, and track reading) to represent the movement of each character in the scene-it may be hand-drawn or computer-generated. Initial hand-drawn ruff animation yields material that typically omits frames, making fine-timing adjustments in the scene

Presented at the 140th SMPTE Technical Conference (paper no. 140-1), in Pasadena, CA, October 28-31, 1998. Leonard J. Reder and Gene Takahashi were with Warner Bros. Feature Animation, Glendale, CA at the time the paper was presented. An unedited version of appears in DTV Reality Check-The First Year, SMPTE, 1998.

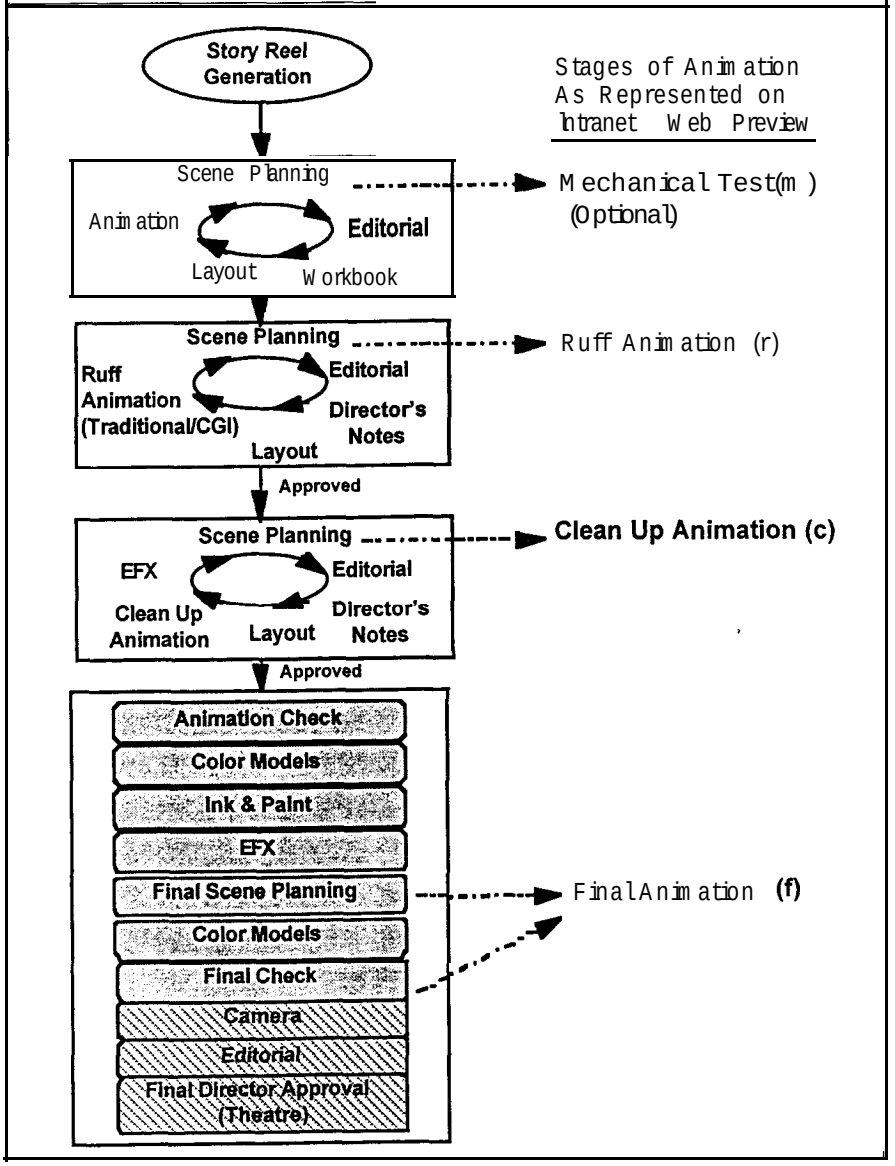

Figure 1, Animation process flow and iteration.

easier and less work for the artist. Ruff computer-generated animation has all the frames but is rendered as wireframe models.

After ruff animation is completed the material is scanned into the computer and goes to the scene planning department that assembles the ruff animated elements, adds the camera moves, and then composites them. Once this process is completed, the sequence of composited frames is sent electronically to editorial for director's approval or director's notes for changes. This is the first point in the animation process where the scenes are automatically posted to the web. Each scene posted at this stage is labeled on the web page as "Ruff Animation (r)."

Next, the ruff animation moves into a cleanup process. During this phase, line quality of the images is improved, and frames missing from the ruff animation are drawn. At this point, every frame within the scene is animated. This is done for each element of the scene and the material is once again sent back to scene planning for compositing and additional fine-tuning of camera moves. Scene-planned frames are, once again, sent to editorial. The Intranet web software picks up the material at this point and updates the ruff version to cleanup. The scene label on the web page changes from "Ruff Animation (r)" to "Clean Up Animation (c)." Again, the director's approval or notes are obtained before the scene can continue on to the final stages.

Once clean up animation is approved, the scene is ready to have final color added. This part of the process involves the greatest number of animation departments: first, the animation check department carefully examines the animation and fixes minor flaws or sends the scene back to a previous department to be fixed. After approval, the scene goes to the color models department where a palette that determines how the scene will be colored is defined; if approved by the art director the scene is inked and painted (colored by computer). Next, the completed animation goes through digital effects, final scene planning, and color models departments for further refinement. Final check does a meticulous and thorough inspection of the frames of animation before anything is recorded onto film. During this process the web pages are updated. The scene label changes for the last time to "Final Animation (f)." The final colored animation is recorded to film using Celco film recorders and viewed in the screening room for final director approval or notes. Approved film is then cut into the film reels. By this time the approved scene material is on the web and cut into the Avid nonlinear system to keep electronic cuts current with the physical film reels.

The above process can be iterated at several places and for some time depending on the complexity of the scene being animated. Typically, changes occur during ruff or clean up animation with fine adjustments taking place during the final animation phase. Figure 2 is an example of how a frame for scene preview, when served, looks at various stages of the process. Within Fig. 2 frames (a), (b), and (c) are the dragon from Quest for Camelot while frames (d), (e) and (f) are the Hogarth character from The Iron Giant.

Editorial Server Road Map

Within Warner Bros. Feature Animation, several Intranet websites have previously been set up to provide various organizational information; however, none of these served production audio or QuickTime movies. Since we anticipated that media would be streaming from our system at some point, we decided to set up a separate

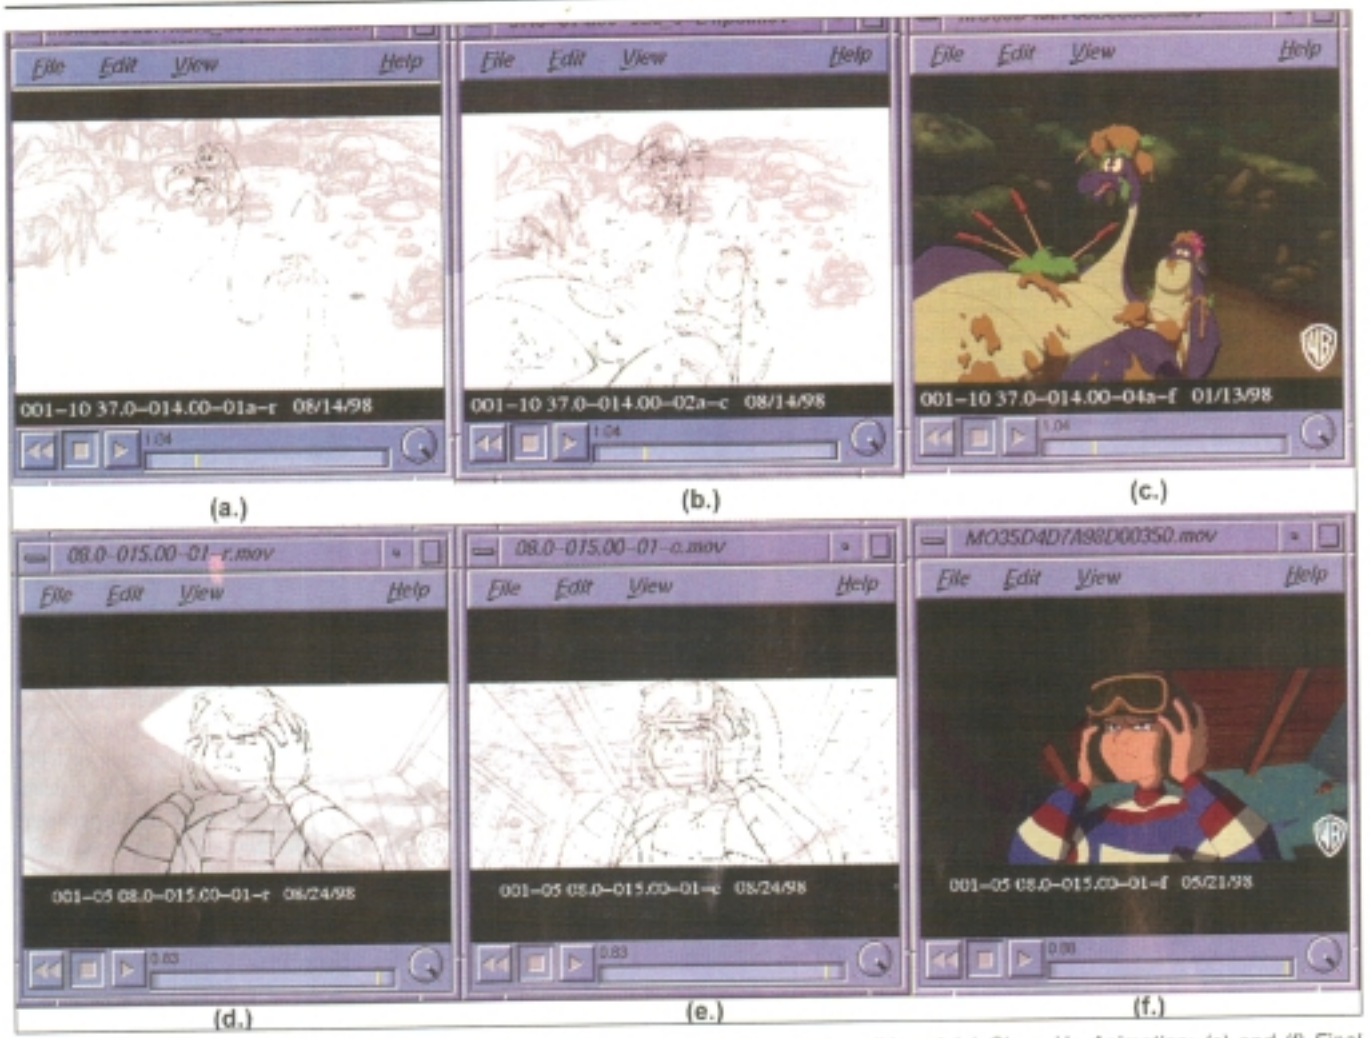

Figure 2. Frame of QuickTime movie for scene preview: (a) and (d) Ruff Animation; (b) and (e) Clean Up Animation; (c) and (f) Final Animation. Top row is frame from Quest for Camelot, bottom row is frame from The Iron Giant.

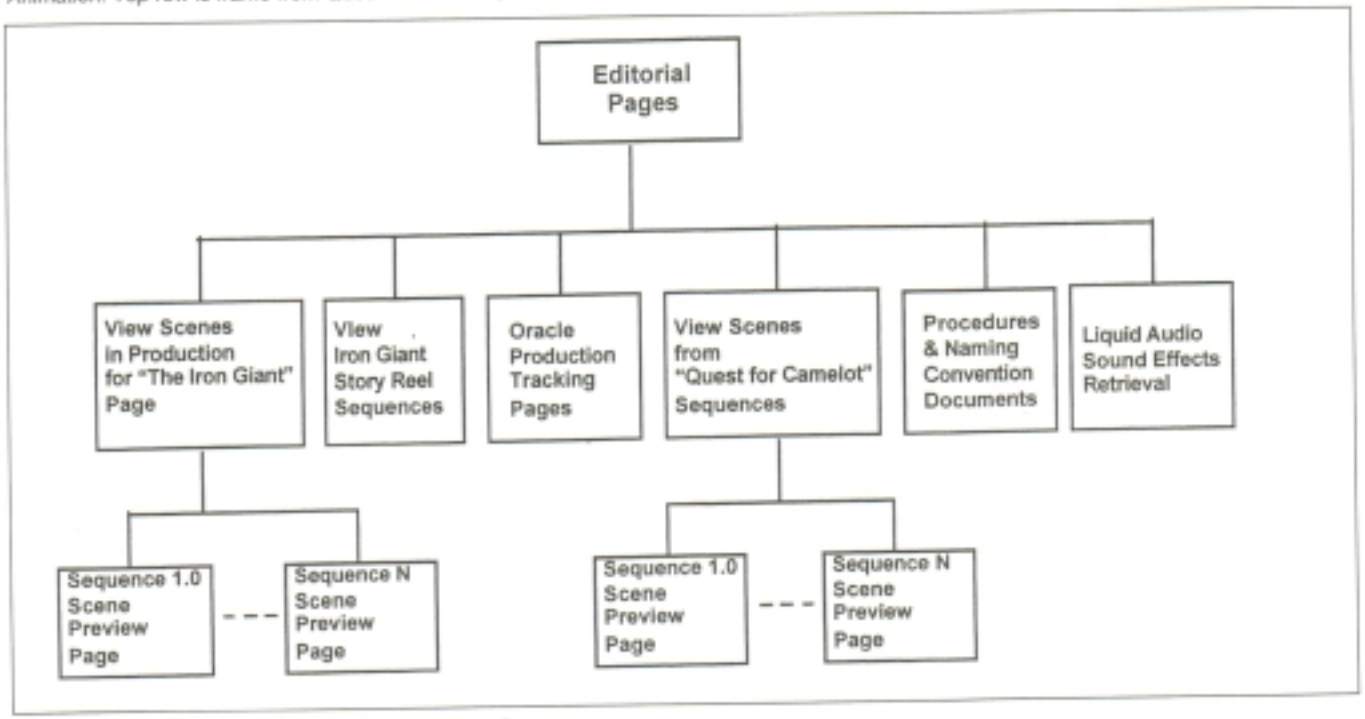

Figure 3. Intranet editorial.wbfa.com Web site tree diagram.

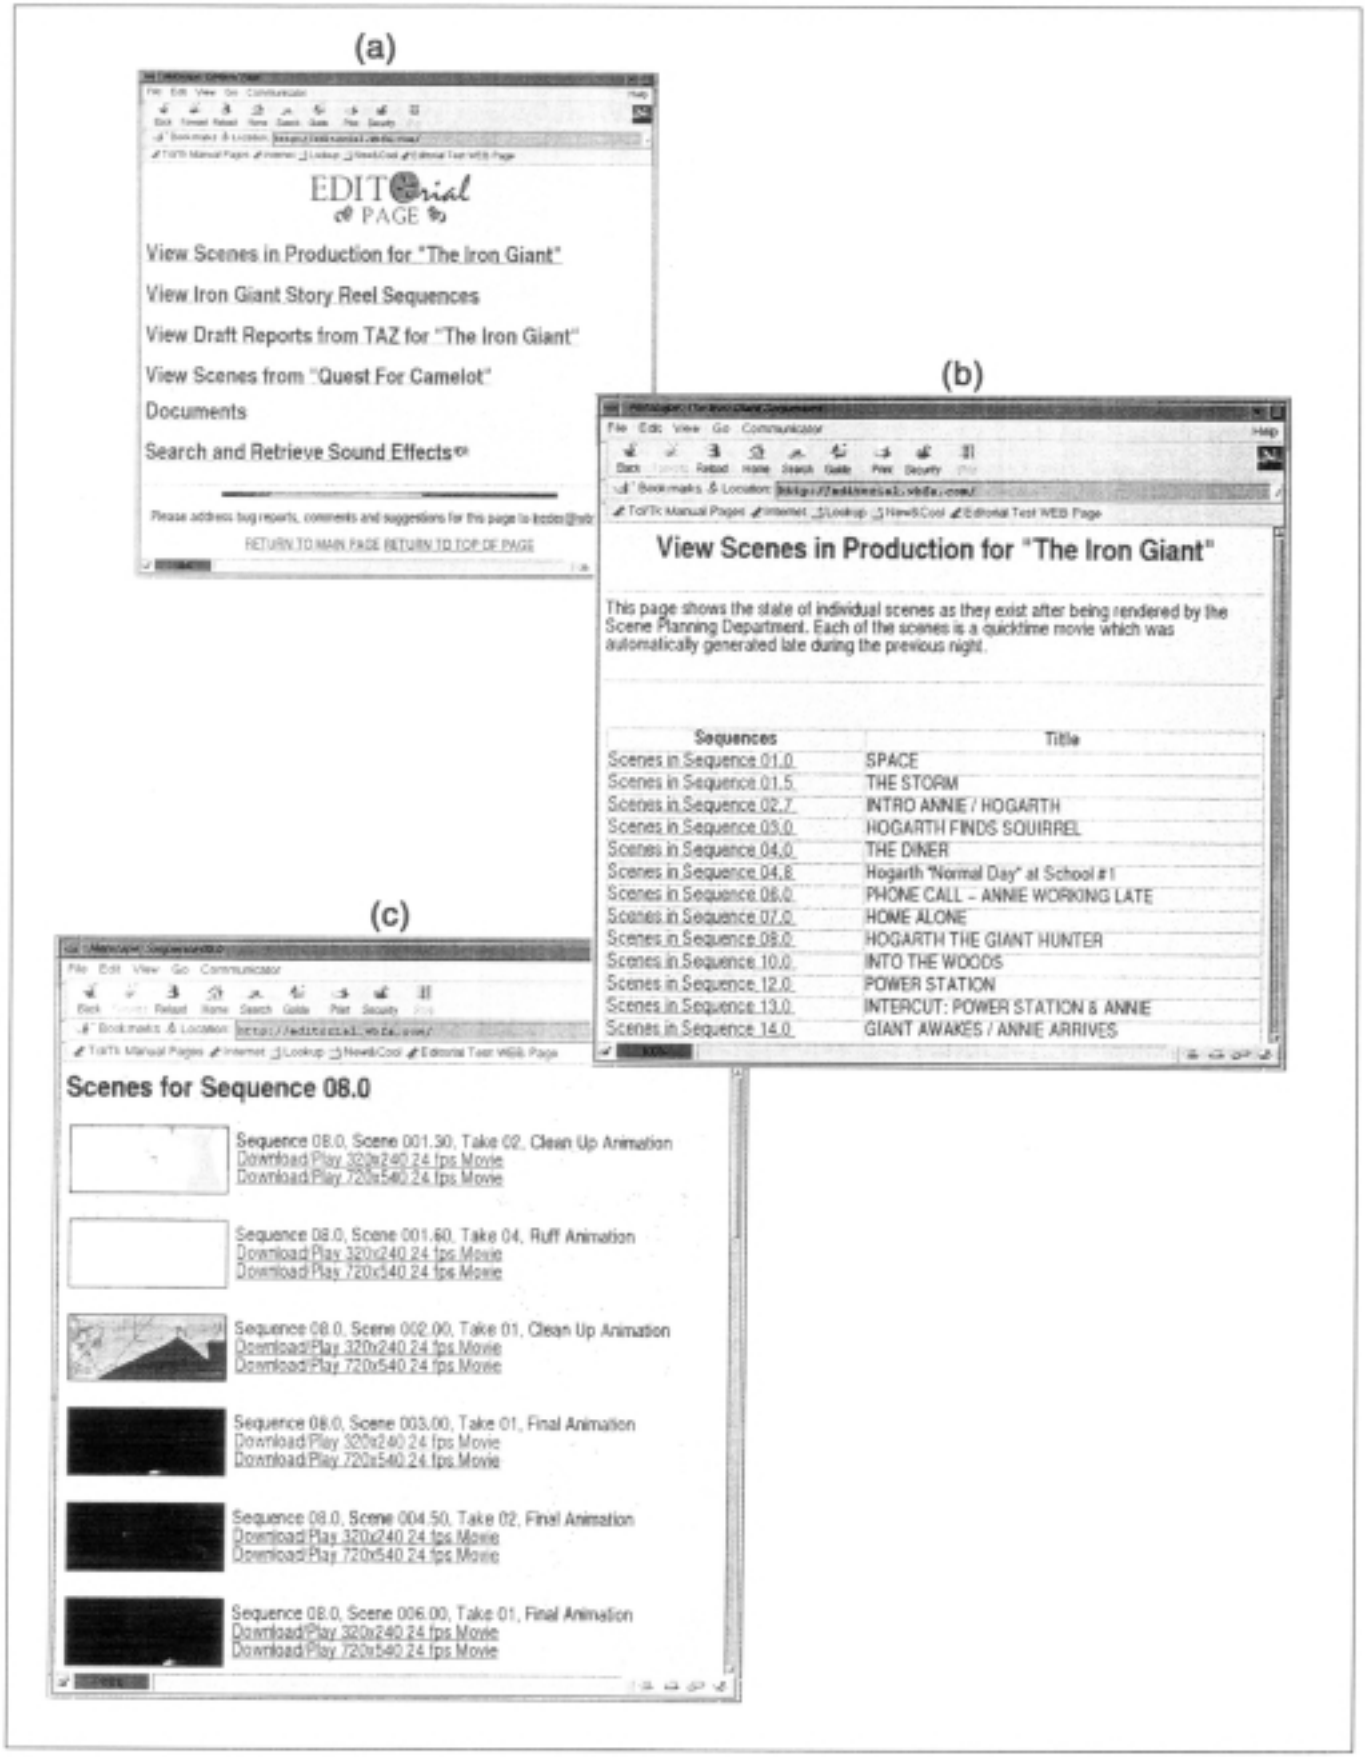

Figure 4. Editorial page progression: (a) editorial page; (b) view scenes in production for The Iron Giant; (c) scene for Sequence 08.0.

web server. The setup is completely production driven in that with each new production the Intranet site, configuration is modified and enhanced to suit the needs of the production. To date, the editorial site has been used to serve material for Quest for Camelot and The Iron Giant productions. Quest for Camelot was released May 1998 and The Iron Giant is currently in production.*

Figure 3 shows the top down organization of the editorial web pages. Web pages utilized for preview of scenes are updated and then frozen to preserve reference material online. The animation departments use this reference material since high-resolution material for the completed production has been archived to tape.

At the top of the diagram is the "Editorial Page" connected by six hypertext links to other pages. A hypertext (or hypergraphic) link is text (or a graphic image) one can click on to retrieve new data associated with that link.

Figure 4 shows the progression of web pages to preview a scene. Starting at the page shown in Fig. 4(a), and clicking on the "View Scenes in Production for The Iron Giant" link takes us to the page shown in Fig. 4(b). The "View Scenes in Production for The Iron Giant" page is generated daily and lists all the sequences that are currently being animated for the film. Next, clicking on a sequence number, for example, Sequence 08.0, takes the user to the page shown in Fig. 4(c). The "Scenes for Sequence 08.0" page presents a list of scenes that can be viewed. Pages with the same format as this page are generated and updated for each sequence of the movie in production. These are called scene preview pages because they are used to obtain the QuickTime movie for viewing any particular scene within a production. As the movie production progresses the pages grow and eventually every scene in the film is served.

Figure 5 shows an enlarged view of the top and bottom of the "Scenes for Sequence 08.0" page. The bottom indicates that the scene order was specified by the editorial cut list and that the page was last updated on July 30, 1998.

Once the user is in a scene preview page for a particular sequence (Fig. 5) the operation is very simple, since each scene detected has a record on the page. The record consists of a small 205 x 87 pixel thumbnail image of the first frame of each scene. To the right, a label is created consisting of the sequence, scene, take numbers, and the stage of animation being presented. Below this are two "Download/Play" hypertext links to the QuickTime material generated for the scene. The thumbnail image is also a hypergraphic link to the low-resolution 320 x 240 QuickTime movie. When any of the three links are clicked on, the QuickTime movie for that scene is downloaded and played on the native movie player of the local machine.

Returning to Fig. 4(a), the next link is "View Iron Giant Story Reel Sequences." This page is automatically generated with links to low-resolution story reel QuickTime movies that the editorial department has uploaded from their nonlinear edit systems. These movies are uploaded after initial approval out of the workbook process and only updated on an individual basis if substantial changes have been made to them.

The remaining four links (Fig. 4(a)) are a top level report page to our Oracle-based production tracking system (called TAZ); scene previews for Quest for Camelot; a page referencing documentation on procedures and naming conventions; and the Liquid Audio Sound Effects Retrieval system.

The TAZ reports are served elsewhere but the Liquid Audio Sound Effects are served as part of the editorial server. This retrieval capability allows the user to search a database containing several sound libraries and to "stream" Dolby AC2-encoded sound for auditioning. A compact disk-quality audio file can then be downloaded for use in either a Mac or Windows environment.

Finally, the "View Scenes from Quest for Camelot" hypertext link points to a page which lists all sequences within the Quest for Camelot production. Once again each sequence is a hypertext link to a page containing a list of scene preview movies.

Production Directory Hierarchy

The scene preview pages described above and associated QuickTime movies are generated based on naming conventions utilized within the production directory hierarchy. The production directory structure begins at a top-level directory. Directly under it are separate subdirectories representing material for each production. Each production is represented with a production descriptor named sub-directory. Under the production directory, the structure is further broken down into a directory for each sequence in the motion picture and under each sequence directory scene directories are created. Under each scene directory are various other subdirectories that contain other elements of data required to assemble and color the animation.

Also, under the scene's sub-directory is a directory containing the completed animation for each of the stages shown in Fig. 1. Directly under this, each version of rendered animation is stored in another directory. The directory structure in UNIX notation is:

/top-level/<Production Desc>/<Sequence Number>/<Scene Number>/<Intermediate Directory>/<TestId-Take_TypeLDesc>

(e.g. /top-level/XXXX/08.0/024.60/out/001_02_r)

Each directory within the < and > represents a specific naming convention explained in Table 1. Table 1 is the top down naming conventions used for the directory structure.

The software to generate the scene preview pages and QuickTime movie of each scene scans the production directory structure for valid media to be included in the pages. When new scene material is found online, QuickTime movies are generated and the HTML web page is updated with an additional scene record and stored in a special editorial directory located under each sequence directory. We also store production audio within the editorial directories.

Within each editorial directory the HTML page is named SequenceXXX.html where XXX is the sequence number and the QuickTime movies are represented with the naming

* The animated feature The Iron Giant opened in theaters in August 1999.

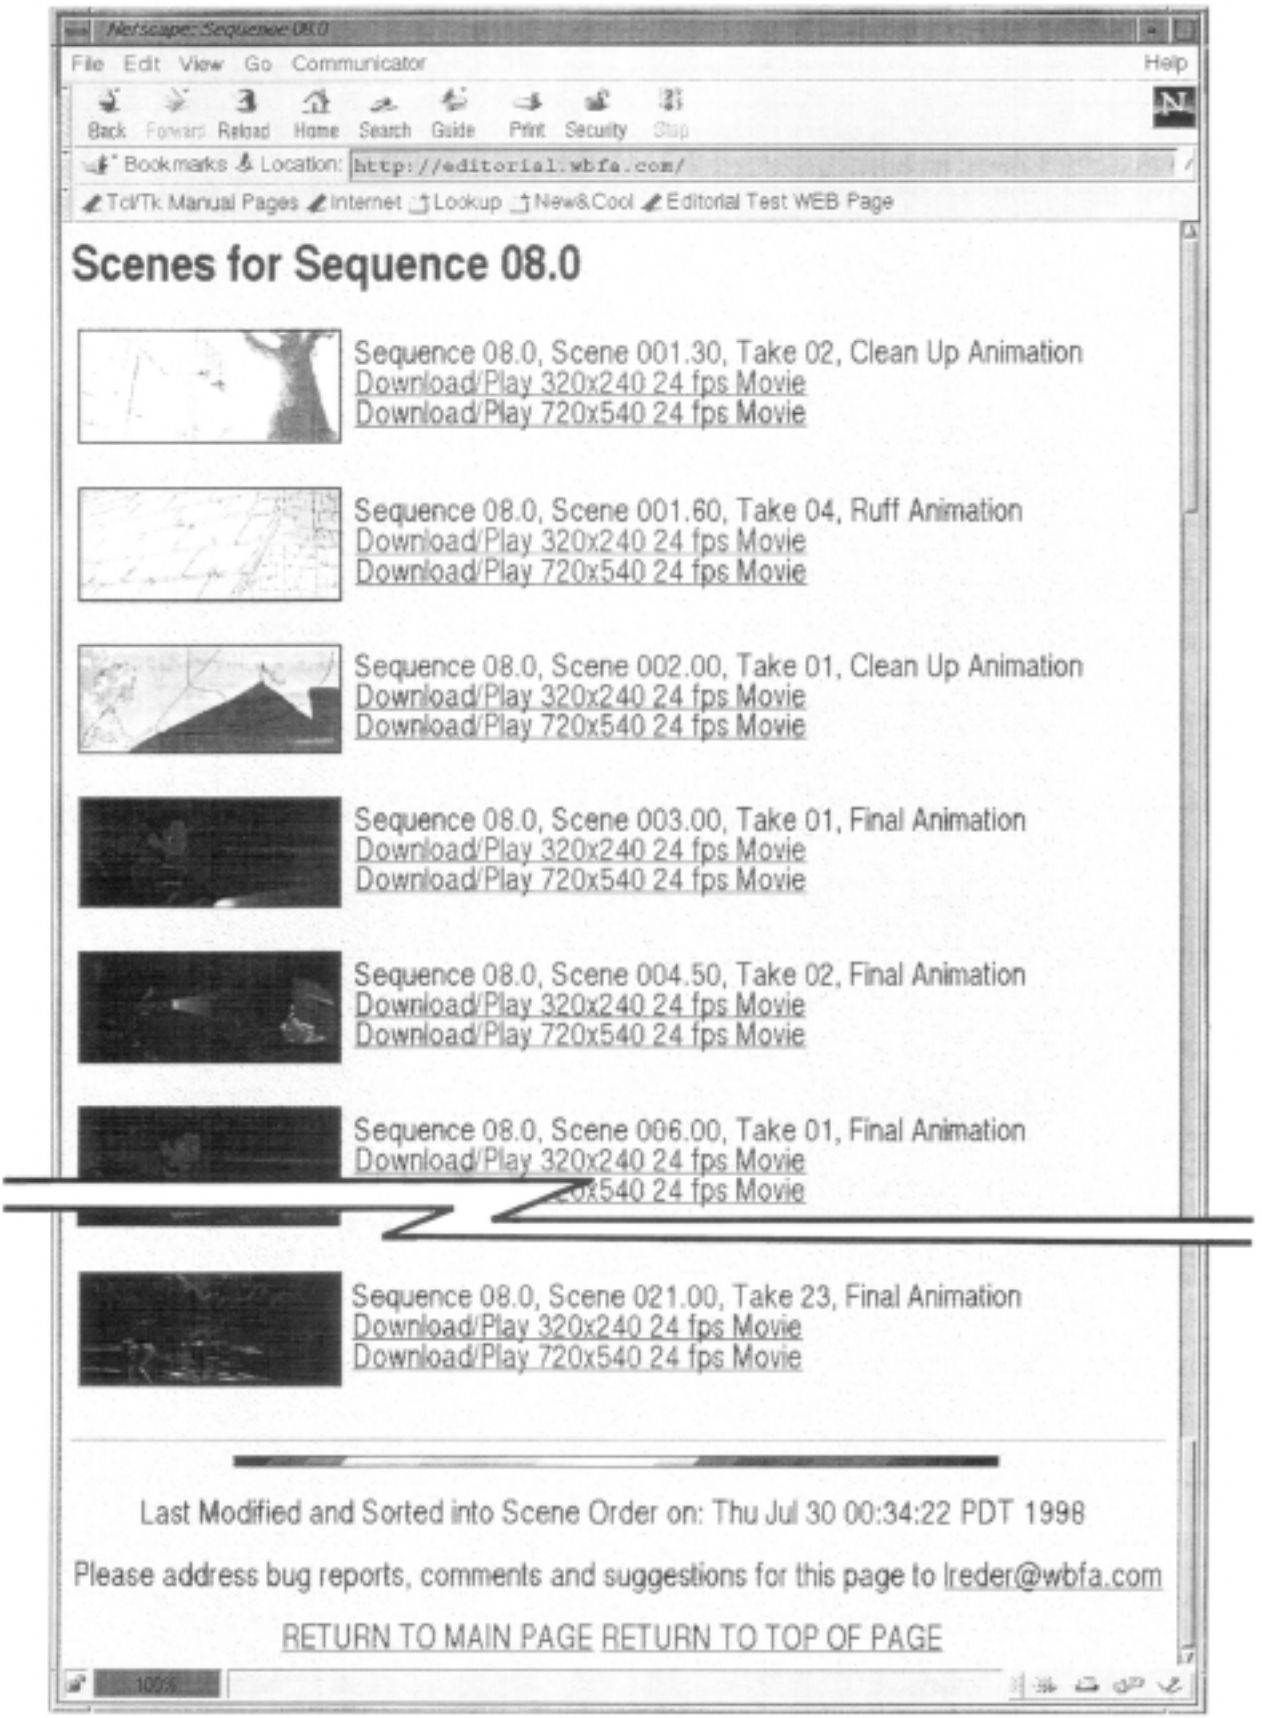

Figure 5. Scene preview page generated for Sequence 08.0 of The Iron Giant production.

vention of:

<Sequence Number><Scene Number><Take Number>-<Type><Movie Type>.mov

(e.g. 08.0-024.60-02.r.mov or 08.0-024.60-02-r-dl.mov)

Each component within the <> of the QuickTime name is explained in Table 2.

Implementation

Initially it was thought that a shareware NextStep application displaying multiple images on the desktop could be adapted to play small sequences of frames. At that time most of Warner Bros. Feature Animation was utilizing 2-D animation software running on the NextStep operating system to digitally manipulate and assemble the animation. The exception was the CGI group running 3-D animation software on Silicon Graphics Inc. (SGI) workstations. The studio is phasing out the use of NextStep in favor of other 2-D software packages; therefore, we considered using the SGI showcase application to reference multiple images and audio for playback on the SGI. The problems with these approaches were:

- They would have required separate implementations on each machine.
- Macintosh users (most of our production management team) would not have the capability to preview.
- We required a workable solution, within a few weeks of concept.
- The solutions would only make a limited number of scenes available to the user at any particular time.

Utilizing popular web and QuickTime technologies would make all animated scenes available for viewing on any type of computer connected to the Warner Bros. Feature Animation Intra-network.

To keep development time short, all software was written utilizing the Perl language. Perl is an advanced scripting language, combining the system Control functions of standard UNIX shell script languages with the ease of string manipulation and file input/output (I/O) found in compiled languages such as C or C++. Moreover, Perl was designed for generating nicely formatted data reports and performing system functions like execution of image manipulation command line utilities or standard UNIX commands.2,3

Table 1—Subdirectory Naming Convention for Production Directory Structure

Subdirectory	Description
top-level	The top level root directory.
Production Desc.	Production descriptor name of production directory (XXXX).
Sequence Number	Sequence directory name (e.g. 27.3 or 08.0).
Scene Number	Scene directory name (e.g. 01.0 or 37.42).
Intermediate Directory	Storage location for various versions of scenes.
Testid_Take_Type[_Desc]	Version of scene rendered directory name. This consists of a three digit test identification number (the Testid) unique to each department and a two digit take number (the Take) which corresponds to director's notes. A type identifier (Type) which is either, m for mechanical test, r for ruff animation, c for clean up animation, or f for final color animation. The _Desc field is optional and typically used for some form of department abbreviation such as _16bit (tells us the scene was rendered as 16bit TIFF) or _fsp (tells us that finel scene planning rendered this version). The [] around the -Desc is UNIX notation for an optional argument.

Table 2—QuickTime of a Scene File Naming Conventions

Name	Description
Sequence Number	Sequence number where scene came from. This is derived from the production directory path.
Scene Number	Scene number within sequence. This is derived from the production directory path.
Take Number	Take number (the Take) which corresponds to directors notes and is also derived From the production directory path.
Type	Type identifier that is either, m for mechanical test, r for ruff animation, c for clean up animation, or f for final color animation. This is also derived from the production directory path.
Movie Type	This is the type of movie generated for page reference. If this is not defined the movie is a 320x240 pixel format 24 fps movie. If this is "-dl" the movie is a 720x540 pixel Dl video format compatible movie. The movie type is generated internally within our soffware.

Three scripts and one Perl library were developed to refresh the HTML web pages and generate the required QuickTime movie. The main script of interest is UpdateScene.pl that generates all the scene preview pages. It is built on the Perl library script named SceneGenlib.pl that handles all image file manipulation tasks and HTML generation. Two other scripts called UpdateSeq.pl and UpdateStory.sh refresh the "View Scenes in Production for The Iron Giant" page (Fig. 4(b)) and the "View Iron Giant Story Reel Sequences" page.

The UpdateScene.pl script runs every night at midnight utilizing the UNIX cron facility on the SGI server. The cron is a program for timed execution of tasks.

HTML

Our pages are entirely implemented utilizing HTML, because it is simple to generate using a string-based scripting language like Perl. The fundamental function of the UpdateScene.pl, UpdateSeq.pl and UpdateStory.sh scripts is to generate and update the HTML pages. A brief discussion on the nature and structure of a HTML document (file) follows to give the reader a foundation to understanding how our web pages function.

HTML is a cross-platform, document-viewing language with links (hypertext or hypergraphics). A hyperlink is a reference point (link) in a document to another location in the same document or some other point in another document. Browsers usually display hyperlinks in a distinguishing way, for example, in a different color, font, or style.

HTML uses American Standard Code for Information Interchange (ASCII) tags. A tag consists of text

enclosed by "<>". The tags are typically used in pairs surrounding text to define formatting or a hypertext or hypergraphic link. For example the word bold expressed as <B>bold</B> would appear as bold typeface within an HTML-compliant browser. Although most common tags appear in pairs, some do not; exceptions are the <IMG>, <P>, <HR>,
 and <! "string"> tags. The common tags utilized within the web pages generated by our software are explained in Table 3.

A hypertext (or hypergraphic) link is implemented with an anchor tag <A> that can enclose text or an inline image. The syntax for the anchor is as follows:

```
<A HREF="URL"> Text to be linked</A>
```

Text that has an anchor around it appears underlined and in a different color. If an image is a hypergraphic, then it has a colored border surrounding it. In our implementation, this is the means by which QuickTime movies are acquired onto a local machine for playback.

The thumbnail images that appear on the scene preview pages (Fig. 5) are embedded inline with the HTML tag <IMG SRC="URL">. The URL's (Universal Resource Locator) syntax is "protocol://Internet-address/file name." The "protocol" is typically hypertext transfer protocol (http) but file transfer protocol (ftp) is sometimes used. The protocol can also be specified as "file" meaning the browser should access the local file system of the machine. The "Internet-address" is the name of a server on the network. As a shortcut, the "protocol" and "Internet-address" can be left off the URL. When this is done, the default is to use http as the transfer protocol and the original server hosting the page address as the "Internet-address."

The structure of an HTML page consists of two parts, the header and the body. The header defines meta-data (data about the document itself), while the body is where all the formatting tags that define the look of the document are placed. The structure is as follows:

```
<HTML>
  <HEADER>
    Title of page appears as title of browser
    window. Other data about this document is also
    placed here.
  </HEADER>
  <BODY>
    Text, images, formatting tags and hypertext
    (or hypergraphic) links within the document.
  </BODY>
</HTML>
```

In the course of implementing the web pages for this site, several characteristics of HTML became apparent. Text within an HTML document not surrounded by formatting tags will word wrap to the current width of the browser. Within the HTML document any extraneous white spaces, tabs, or carriage returns are ignored. This means that if specific formatting within a page is desired, it must be done

Table 3 —HTML Tags used in EditorialWeb Pages[5]

Tag Name*	Description
<IMG SRC="URL">	The IMG tag specifies an image file to be displayed inline within the document. The SRC="URL" attribute is mandatory to specify the image file URL. Other attributes can be given to alter the displayed image's width, height, alignment, etc.
<P>	Marks beginning of a new paragraph (or logical block of text).
<HR>	The HR tag causes a horizontal line to be drawn across the width of the browser's document viewing window. HR also can have a number of attributes to change the line color, width, etc.
 	Character which causes a line break to appear
<! "comment string">	The <! is the start of a comment string and the > is the end. The "comment string" will not appear inside the browser's document window but will appear in the HTML source.
<HTML>...</HTML>	This is the outer most tag and identifies the document as an HTML file. HTML is a subset of the standard generalized markup language (SGML) that is an ISO standard.
<HEAD>...</HEAD>	The HEAD tag defines the meta-information for an HTML document file. A set of tags is used within the HEAD tag set to define several properties of the document. It is mandatory that a TITLE tag be specified within the HEAD tag set.
<TITLE>...</TITLE>	Text within the TITLE tags defines the title of the document. This title does not appear within the document but rather is displayed as the window title of the browser frame that the document is being displayed in.
<BODY>...</BODY>	The BODY tags define the body of the document. All visible text within the body is than enclosed within other tags to give formatting and hyperlink information to the document. All tags that define the look of the document are contained within the BODY tags. The BODY tag can have several attributes to give the document a particular look. In Listing 1 the BACKGROUND= attribute defines the look of the background of Figure 5 by causing the paper.gif image to be tiled over the entire background within the browser.
<A HREF = "URL">...</A>	The A tag is an anchor which defines a hypertext link to a URL using the HREF attribute. Text or graphics within the A tag is typically highlighted in some way. Other attributes are also defined for anchors but at least a HREF or NAME attribute must be used.
<Hn> </Hn>	The tags Hn (where n = 1,2,3....6) represent headings. Although browsers do not associate the heading with a particular font, the convention is to use H1 for the largest or most import heading and assume "n" goes up as the font size progresses downward. By default headings are left justified but an optional ALIGN attribute can be specified to over ride this.
<B>...</B>	Text enclosed by the B tags appears as bold.
<I>...</I>	Text enclosed by the I tags appear as italic.
<TABLE>...</TABLE>	Used to define a table within a document. It is used within Figure 5 to cause text to be centered along side of an image. In Figure 4b a standard table is created. Several attributes are used to define the look of the table. BORDER, WIDTH, CELLSPACING, CELLPADDING are some of the common attributes used to define tables. BORDER=number causes a line of thickness number to be drawn around each cell of the table. If BORDER=0 is used a borderless table is created. WIDTH specifies the table width is either absolute units of pixels or in percentage of window width. CELLSPACING=number specifies the spacing, in pixels, to be placed around text or images within a table. CELLPADDING=number specifies the space, pixels, between individual cells in a table.
<TR>...</TR>	The TR tag defines a row within the table. It is used within a TABLE tag.
<TD>...</TD>	The TD tag defines a cell within the table. Actually the TD tag defines a data cell within the table and the TH tag defines a header cell within a table. Both of these are used within the TR tag.
<CENTER>...</CENTER>	The text enclosed by the CENTER tags is center aligned. A line break is added before and after the enclosed text so only the text to be centered is on the line.
<FONT>...</FONT>	Text is enclosed by the FONT tag to control size and color properties using the SIZE and COLOR attributes of the FONT tag. The SIZE=number specifies the font size as an absolute range from 1 (the smallest) to 7 (the largest). The COLOR=#rrggbb specifies the fonts RGB color using a triple hex number representation.

*HTML tags that do not require pairs

with HTML formatting tags. It is important to realize that all HTML formatting features "specify logical rather than physical" characteristics of a document's look. This means tagging some text as a header (using the <H1> and </H1> tags) will, on a Macintosh, map to some font selected by the local browser while on an SGI workstation it may map to a completely different font. While identical pages can appear significantly different on various systems the page's functionality does not change. We have designed our pages with a consistent look across all systems."

Custom Perl Software

Within the UpdateScene.pl script, we update scene preview pages, consisting of a header, repeated instances of a scene record, which points to the movies to be viewed, and a tail section. Listings 1, 2, and 3 show the HTML of the page shown in Fig. 5, and Table 3 explains each of the HTML tags used.5 The code generates header and tail HTML for the top and bottom of every page. In between, individual records for each scene to be previewed are generated using the format shown in Listing 2, which starts with a scene record HTML comment line. These comments do not appear to the user viewing the page but are used by the software for sorting the scene record in top to bottom order according to editorial cut lists stored on the central file server.

Figure 6 shows the processing flow within the UpdateScene.pl script for creating and updating material. The script takes a production number as its argument and then iterates through the sequence and scene directories. For each valid scene output detected, a list of possible scene versions is generated within the code. This list is then searched for the most current version of the scene. First, the type identifier is examined for f, c, r, or m with the f for final type having highest priority, c for clean up next highest, and so on. Only scene versions with the highest priority type are retained. Next, only the scenes with the highest take number are kept, and finally the scene with maximum test identification is used. This process is performed with the searchforscene Perl function, and utilizing the scene path returned, the update is performed by the UpdateSceneMovie Perl function.

Within UpdateSceneMovie, the routine CheckAndUpdateQt is used to check for existence of a rendered QuickTime movie, a GIF thumbnail of the first frame of a scene, and the .QtSceneRecord-XXXXX file containing information about the last scene update made. If all these exist, then processing continues to the CompareAndUpdateQt routine. However, if any of the above are missing, then either a new scene is online or something has been inadvertently deleted and a new QuickTime movie and HTML update are performed. At this point the script iterates to the next scene.

If CheckAndUpdateQt returned and indicated the scene existed normally, then the CompareAndUpdateQt routine is called. The routine compares the scene found by searchforscene with the last version of the same scene that exists on the web page. If the scene is newer, the QuickTime and HTML are updated.

A simple distributed database is used to track the state of each scene on every web page. The UpdateScene.pl script reads and writes out a scene record file to and from the appropriate sequence editorial directory. The file is named .QtSceneRecord-XXX.XX (for example .QtSceneRecord-024.60) with XXXXX representing each scene number. These are hidden UNIX files. Within the .QtSceneRecord-XXX.XX

Listing 1. HTML from Top of Figure 5

Header HTML Source for Sequence 8.0 Scene Preview Web Page

```
<HTML>
<HEAD>
<! DO NOT EDIT THIS FILE - IT IS MACHINE GENERATED >
<TITLE>Sequence 080</TITLE>
</HEAD>
<BODY background="/top-level/XXXX/editorial/www/images/paper.gif">
<H1>Scenes for Sequence 080</H1>
<P>
<! START OF SCENES >
```

Listing 2. HTML Record Used in Figure 5

Scene Record HTML Source of Sequence 8.0 Scene 4.5

```
<! XXXX-080-00450-014-02-f>
<TABLE BORDER=0 CELLSPACING=0 CELLPADDING=5>
  <TR>
   <TD ALIGN=center>
      <A HREF="/top-level/XXXX/080/editorial/qt/080-00450-02-fmov">
     <IMG SRC="/top-level/XXXX/080/editorial/qt/080-00450-02-fgif><br>
      <A>
   </TD>
   <TD ALIGN=left>
      Sequence 080, Scene 00450, Take 02, Final Animation <br>
     <A HREF="/top-level/XXXX/080/editorial/qt/080-00450-02-fgif"><br>
      Download/Play 320x240 24 fps Movie<br>
      <A>
    <A HREF="/top-level/XXXX/080/editorial/qt/080-00450-02-fmov/">
      Download/Play 720x540 24 fps Movie<br>
      <A>
   </TD>
  </TR>
</TABLE>
<P>
```

Listing 3. HTML from Bottom of Figure 5

Tail HTML Source for Sequence 8.0 Scene Preview Web Page

```
<! END OF SCENES >
<P>
 <center>
  <P><HR><IMG SRC="/top-level/XXXX/editorial/www/images/linespectranice.gif"><P>
  Last Modified and Sorted into Scene Order on: Thu Jul 30 00:34:22 PDT 1998
<P>
  <font size=3
   Please address bug reports, comments and suggestions for this page to
   <A HREF="leslie@wbfa.com">leslie@wbfa.com</A><BR><BR>
    <A HREF="http://editorial.wbfa.com/">RETURN TO MAIN PAGE</A>
    <A HREF="/top-level/XXXX/080/editorial/qtSequence080.html">RETURN TO TOP OF PAGE</A>
  </font
 </center>
<BODY>
</HTML>
```

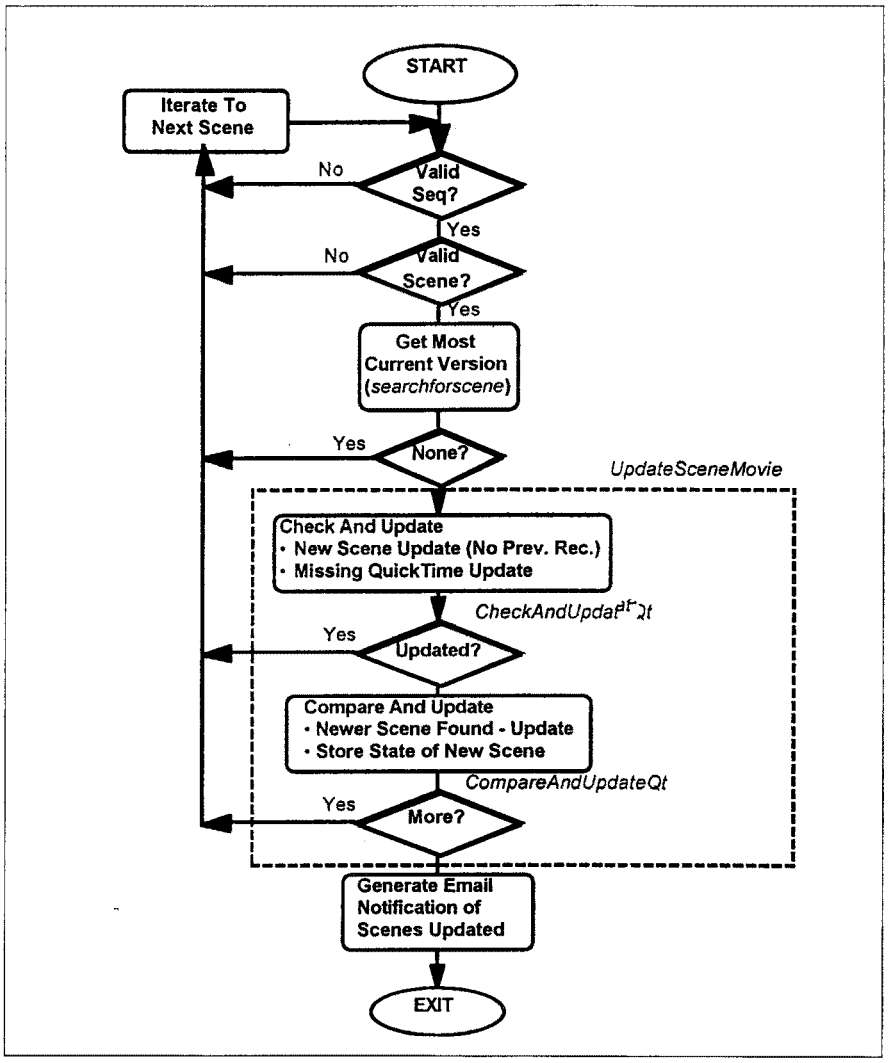

Figure 6. Structure and logical flow of UpdateScene.pl script used to regenerate HTML pages and QuickTime movies for scene preview.

Table 4 — Interface Functions of SceneGenlib.pl

Function Name	Description
RmMovGif	This routine removes the old instance of the GIF image file and QuickTime movie file of the scene last displayed.
MakeThumbnailGif	This routine generates a thumbnail GIF image file of the first frame of the scene being updated or created .
MakeAnnoTiff	This routine generates a temporary directory filled with annotated TIFF images. These images are used by MakeQtSceneAnnotated to generate the QuickTime movie of each scene.
MakeQtScene	This routine generates the QuickTime movie file without annotation for the scene being updated or created.
MakeQtSceneAnnotated	This routine generates the annotated QuickTime movie file for the scene being updated or created.
MakeQtSceneWithAudio	This routine generates a QuickTime movie with audio from an input silent QuickTime movie and a input AIFF audio file.
UpdateHTMLSceneRecord	This routine checks to see if an HTML page for a sequence exists. If not than a new HTML page with standard header (Listing 1) and standard tail (Listing 3) template text is generated and the HTML scene record of Listing 2 is added. If the HTML exists than either add or update the HTML scene record.

file we store the scene record HTML comment, total number of frames of the scene, time and date the scene was last updated, and the names of the QuickTime movie files associated with the scene. Most of this information is used for debugging purposes, but the scene record HTML comment is used for determining if a scene version has changed. A list of all updated scenes is generated during this process. After all directories have been scanned and processing is completed, an e-mail message is generated to notify users of exactly which scenes have changed on the web pages.

Perl functions contained within the *SceneGenlib.pl* library script take care of the mechanics of generating the QuickTime and HTML (Table 4). The *UpdateHTMLSceneRecord* generates and updates all HTML text produced by the scripts using native Perl functionality. The remaining functions of Table 4 utilize binary executable command line programs, listed in Tables 5 and 6 to perform the required image processing needed to produce the QuickTime movies and thumbnail images. The SGI utilities listed in Table 5 are standard with any model of SGI workstation.[6] The utilities in Table 6 are public domain software available from the San Diego Supercomputer Center (SDSC).[7]

The original production materials are sequences of frames stored as either 1024 wide or higher film-resolution TIFF image files. The files are either 8 bit or 16 bit per component RGBA format depending on the scene. The height of the images depends on the production. Both 1.66 and 2.35 aspect ratio using 35mm 4-perf film have been shot within the studio. *Quest for Camelot* was a 1.66 show and *The Iron Giant* is being shot in 2.35 CinemaScope. Various versions of the software adjust for either format.

To generate the QuickTime of the scene the actual process is as follows. The first step is to resize the TIFF images into 720 x 540 images with letterbox at top and bottom. These images are stored in a temporary directory where we generate a sequence of annotation strip images containing the annotation label for each frame. These annotation strip images are created utilizing a com-

Table 5 — SGI IRIX End-user Image and Media Tools[6]

Program Name	Description
imginfo	This command line application prints various information about any file format supported within the SGI ImageVision Library. Information returned includes the image dimensions in pixels, data type of pixel component, dimension order, color model used, the minimum and maximum pixel values, file format, type of data compression. Our applications utilize this program to obtain the image dimensions.
izoom	This command line application generates a resized image output file from an input file. The output resized image file has been filtered using a 2-D convolution with a triangle filter kernel to enhance the appearance of the resized image. Other filters can be selected.
dmconvert	This command line application can be used to convert many different types of image, audio and movie file formats to many other types. The application can alter characteristics of both input and output audio and/or video. For example this application adds the letterbox automatically to our frames. In addition 24 fps movies can easily be converted to 30 fps. Sample rates of audio can be changed. This application has many command line options and is very versatile.
makemovie	This command line application can generate QuickTime movies from input image and audio data from a variety of different formats. The makemovie is used to combine silent QuickTime movie files and AIFF audio files to add sound to our scene previews.

Table 6 — SDSC Image Tools[7]

Program Name	Description
imcomp	This command line application combines two input image files together using digital compositing operations. The compositing uses one or both of the input image files alpha channel as a coverage mask. In the case of the WB logo it is replacing the original pixel data within the image.
imconv	This command line application converts an input image file using one file format to an output image file with a differing file format. It was used here for its ability to convert 16 bit per component TIFF to 8 bit per component TIFF.
imcopy	This command line application copies a portion of an input image into a new file. Input and output files may have different image file formats.
impaste	This command line application copies an input image onto a background image and saves the result into a new file. Input, background, and output files may have different image file formats.

mand line application called *anno*, which is an X-Windows C program. Next, a small Warner Bros. logo is composited onto every frame. The annotation strip images are overlaid onto the bottom letterbox part of each image.

A QuickTime movie is generated from the preprocessed TIFF files utilizing the SGI *dmconvert* utility. Finally, a check is made for the existence of an AIFF format audio file stored on our server, if found the *makemovie* utility is executed to stripe the 320 x 240 pixel QuickTime movie with production audio. The 720 x 540 QuickTime movies are currently not striped.

Each frame is annotated to be consistent with editorial media used within our Avid nonlinear edit systems and film shot by our camera department. The annotation is shown in Fig. 2 and consists of the foot-frame number of each frame relative to the start of a scene, a string indicating sequence, scene, take number and type id, and the date the movie file was generated.

Every day at noon, a script called *UpdateSeq.pl* is executed. By this time the *UpdateScene.pl* processing will have been completed. When a new scene from a new sequence is placed online the *UpdateSeq.pl* script updates the "View Scenes in the Production for The Iron Giant" page of Fig. 4(b) with the new sequence links. The script also acquires the titles for each sequence from our Oracle production tracking system and places them on the page.

Finally the script named *UpdateStory.sh* runs every two hours to regenerate the "View Iron Giant Story Reel Sequences" page, which is linked to the low-resolution story reel QuickTime movies uploaded to the file server by the editorial department. After the page has been updated e-mail notification indicating which versions of the story reel sequences have changed or have been added is automatically generated and sent to select users.

Client/Server Intranet Configuration

Figure 7 shows the client/server Intranet configuration utilized within Warner Bros. Feature Animation. None of the functionality discussed above would have been possible without the infrastructure previously put in place by the technology department at the studio. Two SGI Challenge XL computers serve the production directory file system. Most SGI workstations are connected to the central file servers via ATM (OC 3 155 Mbits/sec) to the desktop and access the remote production file system utilizing Network File System (NFS) Version 3. The remote file system is accessible to SGI workstations at speeds comparable to the performance of local disk access. The editorial http server accesses the actual QuickTime media using NFS and sends the requested files to client machines using http.

HTTP and MIME

HTML browsers are built utilizing HTTP, which defines the messages sent between the client browser and web server. All communications needed are supported by the HTTP/1.0 protocol that is documented in RFC 1945.[8]

When either an HTML page or a QuickTime movie is opened, the browser sends an HTTP request for data. The server replies with an HTTP message containing several pieces of information. Of particular interest is

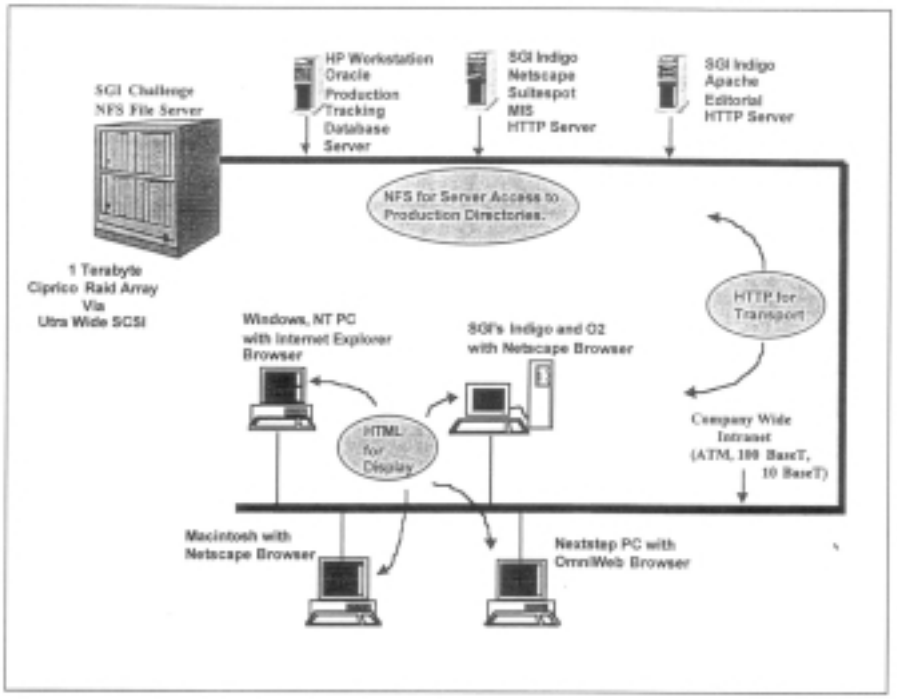

Figure 7. Client/server Intranet configuration.

the message string "Content-type: type/subtype" sent by the server before the transmitted data.

The browser sees this message as an indication of what type of data will be received. The "type/subtype" label is called the Multipart Internet Mail Extension (MIME) type. The type field can be any of the following: image, audio, text, video, application, etc. In our application, we use text/html and video/quicktime. On the server side a file exists (such as, /usr/local/etc/httpd/conf/mime.types) which associates specific MIME types with file name extensions. Most browsers maintain some sort of mime.type file on the local machine, but configuration is done through the browser's user interface. By configuring the MIME type association on the browser, the native machine can launch particular applications based on the type received. For example on our SGI workstations, when Netscape receives a video/quicktime MIME type, the SGI *movieplayer* is launched to play the QuickTime. Recently many software plug-ins have been developed for browsers so that the browser itself can handle the media when the correct MIME type is detected.

HTTP Server

Our web server choice was the Apache web server, which easily accommodates custom software modules. The name Apache was derived from "A PAtCHy server" since it was written from many patches contributed by many developers.[9]

Storage

A given production consists of thousands of scenes. The *UpdateScene.pl* software generates two QuickTime movies for every scene in the film. Because QuickTime is low-resolution and compressed, it requires much less space than the high-resolution production material. Our central server is configured with one terabyte (1000 Gbyte) of disk space. The *Quest for Camelot* scene previews required 28 Gbytes of disk space while *The Iron Giant* is currently utilizing 9.3 Gbytes. Running time for *Quest for Camelot*, less the credits, was 76:42 min and *The Iron Giant* will be about the same. Each film's storage utilization for preview movies is about 3% of the total online storage available. It would have been desirable to retain more than the current version of a scene but doing so would have required excessive disk space.

Security

Security was another issue in the initial implementation stage. Our facility is not directly connected to the Internet but rather to the corporate Time-Warner network located at the Warner Bros. Studios in Burbank, CA. The studio's MIS department maintains an Internet firewall with strict security guidelines to keep outsiders from breaching the network. In addition, all QuickTime movies produced for distribution via the internal web have each frame watermarked with a Warner Bros. logo. Any material transmitted to the Internet can be traced. Logs are kept of all material generated by the *UpdateScene.pl* script and all transactions within our web server.

Table 7 — Download Speed Utilizing http with Various Computer Network Combinations

Machine (CPU)/ Network Connection	Download Speed (Kbytes/sec.)
SGI Impact (R4400 MIPS 250 MHz.) ATM	1966
SGI Indigo XL (R4400 MIPS 200 MHz.) ATM	1501
Windows NT PC (Pentium II 133 MHz.) 100baseT	1179
SGI O2 (R5000 MIPS 180 MHz.) 100baseT	1033
Macintosh Model 9600 (PowerPC 350 MHz.) 10baseT	496
NextStep PC (Pentium II 133 MHz.) 10baseT	373
Macintosh Model 7500 (PowerPC 100 MHz.) 10baseT	154

Table 8 — QuickTime Movie Storage Requirement

QuickTime Format	Storage (Kbyte/frame)	(Kbyte/foot)*
320 x 240 Apple Video Codec	71 - 100	1136 - 1600
720 x 540 JPEG	13 - 58	208 - 928

* 16 frames per foot when using 4 perf 35mm

Performance

The studio consists mainly of SGI workstations, NextStep PCs, and Macintosh computers; a few Windows NT PCs are also in use for test purposes. All the machines are networked: NextStep PCs are connected to 100baseT (100 Mbits/sec) network and Macintoshes to 10baseT (10 Mbits/sec). Typically, ATM-connected SGIs can transfer files at about 5 Mbytes/sec (via NFS), while the transfer rate of the Macintosh is nominally 200 to 300 kbytes/sec (via ftp). Table 7 shows nominal download performance measurements for each system.

We could have gained significant performance when downloading to SGI workstations by using the NFS protocol over our ATM network. To do this the URLs referencing the QuickTime within Listing 2 must specify a protocol of "file" and access a remote file system. Although this would have improved download performance for the SGI workstation users, our experience with NFS software for Macintosh and Windows NT environments has been that the performance would not be an improvement over the use of http downloads. To keep the implementation simple, we utilized only the http protocol for downloads.

The various computers used within Feature Animation can all download and play the scene preview in a reasonable amount of time. Download time is considerably longer on Macintosh machines due to lower CPU performance and less network bandwidth. Story reel QuickTime requires 5 to 15 min to download when viewed on a Macintosh, but is almost immediate when viewed on the SGI.

The Apple video codec was used to compress the majority of QuickTime, since it will play at or near 24 frames/sec on all machines within the studio, without additional hardware accelerators. The other compression method employed was JPEG, which has a greater compression rate and thus required less disk space. Table 8 shows the storage per frame for each type of QuickTime movie generated. The 720 x 540 QuickTime JPEG D-1 video format can only play at speed with machines that have special hardware accelerators. The SGI O2 workstations have built-in hardware image processing which makes them capable of playing this type of media. On the Avid we never actually play the JPEG compressed QuickTime but rather import it into the Film Composer application for editing and playback.

Enhancements

Two enhancements to the above capabilities can be implemented, the use of streaming rather than download and play and the addition of serving automatically assembled sequences. We opted to use download and play because we could not locate software for playing streaming media on the NextStep PCs; and most streaming servers only have players available for Windows and Macintosh operating systems and nothing for the SGI. Streaming movie playback is extremely desirable when playing story reel movies with long download times on Macintosh computers.

One possible solution to streaming QuickTime movies is to convert all our served QuickTime to what Apple Computer calls "Fast Start" files. According to QuickTime terminology a QuickTime file is made up of data structures called "atoms."[10] Movie generation programs usually store the movie resource atom as the last one within the file. A software filter program can be developed to put the movie resource atom first in the QuickTime file causing the movie player applications to start playing before the entire file has downloaded, effectively streaming over the Intranet.[11] With this approach it is not necessary to invest in special purpose software for streaming.

Currently, people view scenes individually after retrieval from the web. Some individuals have used the SGI application *moviemaker* to connect scene preview movies to look at material in continuity. This has been attempted on the Macintosh using desktop applications such as Premiere and VideoShop. Because editorial cut lists exist online for every sequence in the movie, this means that scenes can be automatically assembled into complete low-resolution QuickTime movies. Eventually every sequence of the film will be served in its current state and updated daily, eliminating the need for manual assembly by the user. It will also dramatically decrease the number of videotapes the editorial department outputs on a regular basis.

Conclusion

Cost-effective scene preview capability has been developed utilizing public domain software and a Perl code developed for Warner Bros. Feature Animation. The Perl code contains under 2200 lines and was implemented utilizing SGI and public domain SDSC image tools. The entire scene preview capability was implemented easily and quickly using the existing computer infrastructure set up for Feature Animation.

The simplicity of its functionality has allowed instant access to scenes for use in verifying scene content and checking for "hookup" problems with surrounding scenes. Scenes can be downloaded and played on Macs, Windows, NextStep, and Unix systems by users with limited computer knowledge; a short learning curve on the basic operations of a web browser is all that is required. Changes that occur in the scenes (as many as 30 scenes modified or created in a day)

are easily detected and e-mail alerts production personnel to scenes moving from one stage to the next.

Other beneficial features have resulted from implementation of the software. The amount of videotape material supplied by editorial to other departments has decreased, reducing tape costs and transfer time. D-1-quality material remains available for reference after the original production material has been archived and removed. It also permits access, by remote facilities, to scenes in production. During the production of *Quest for Camelot,* personnel working at remote locations were able to keep abreast of changes by using the Time Warner Corporate Network to access the editorial server.

In the future, other benefits may be realized by using the system. New streaming technologies such as RealVideo by RealNetworks and Netshow by Microsoft promise a faster play cycle while utilizing less disk space on local computers. This will provide better playback of longer more complex movies allowing access to an even wider base of users. The technology is not restrictive and can adapt to match pace with the evolving computer needs of any production. Less expensive PCs running with NT operating systems can also serve QuickTime movies and create HTML documents.

Warner Bros. has used the scene preview capability successfully on two productions. The web capability is production driven and continually enhanced. The ability to view any scene at any point in the production process by any staff member has been realized, increasing efficiency and reducing costs in the production of animated features.

Acknowledgments

The authors would like to thank the following people for contributing to this effort: Bruce Hatakeyama and Steve Wilzbach for design of the basic directory structure; Lem Davis for the initial assignment to implement the scene preview capability and freedom to use this technology; and Arjun Ramamurthy for motivation to submit and write this paper.

References

1. S. Culhane, *Animation from Script to Screen,* St. Martin's Press, New York, 1988.
2. R. L. Schwartz and L. Wall, *Programming Perl,* O'Reilly & Associates, Inc., Sebastopol, 1992.
3. J. K. Ousterhout, "Scripting: Higher Level Programming for the 21st Century," *IEEE,* Mar. 1998 available at http://www.scriptics.com/people/john.ousterhout/scripting.html.
4. D. M. Chandler et al., *Running a Perfect Web Site,* Que Corp., Indianapolis, 1995.
5. I. S. Graham, *HTML Sourcebook, Third Edition: A Complete Guide to HTML 3.2 and HTML Extensions,* John Wiley & Sons, Inc., New York, 1997.
6. SGI online manual pages, Silicon Graphics Inc. Information can be found at http://www.sgi.com/support.
7. San Diego Supercomputer Center Home Page, http://www.sdsc.com.
8. Apache Week Home Page, http://www.apacheweek.com/features/http11.
9. Apache Server Project Home Page, http://www.apache.org.
10. J. D. Murray and W. vanRyper, *Encyclopedia of Graphics File Formats,* O'Reilly & Associates, Inc.: Sebastopol, 1996.
11. "Makin' Multiplatform Movies" available at http://www.apple.com/quicktime/developers/icefloe/dispatch018.html.

THE AUTHORS

Leonard Reder

Gene Takahashi

Leonard Reder is currently with the Jet Propulsion Laboratory California Institute of Technology implementing astronomical interferometers. His interests include realtime image and DSP processing for scientific instrumentation, video and film post-production technologies, and software development techniques.

Reder was formally with Warner Bros. Feature Animation from 1996 to 1999 and was the principal software engineer assigned to the editorial and post-production departments. His main responsibilities included developing infrastructure applications for integrating nonlinear editing systems, film recorders, and the screening room into the production animation process.

Reder is a member of the SMPTE working group on editorial procedures and practices. He earned a MSEE degree from the University of Southern California and a B.S. in electronic engineering from Cal Poly University at San Luis Obispo.

Gene Takahashi works at Pixar Animation Studios and provides technical support for the editorial departments. He also runs his own company, Multimedia Audio Productions, which specializes in audio and music editing. His interests include audio, video, music, and film production as well as writing.

Previously, Takahashi was employed with Warner Bros. Feature Animation where he was involved in the setup and technical operations of the studio.

Takahashi has published technical articles in various trade publications. He holds a B.A. and M.A. degree in music as well as associate degrees in electronics and in the recording arts.